Personal Levels of Assurance (PLOA)

An AT&T White Paper on Assurance

Authored by J. Oliver Glasgow

With contributions from:

Jeff Dodgen
Kennie Kwong
David Chen
Johannes Jaskolski
Kevin Castellow
Dan Druta

Table of Contents

Executive Overview

This white paper deals with innovation in the area of Identity and Federation in the following ways:

- Introduces an approach for determining transaction-based assurance that speaks to "authentication" based on the necessary current conditions of specified, validated attributes and agreements in a profile.
- As part of these agreements, introduces an ability to provide one-time modular Terms and Conditions that can be consistently applied horizontally across many products and services.
- Introduces technology that allows for anytime revocation of attributes and agreements resulting in only the loss of identified and supported atomic functionality within products and services, and not necessarily the wholesale loss of use of those products and services altogether.
- Introduces an approach for determining the lifecycle of LOA settings for an individual based on the current condition of all attribute declarations whether they are validated or not, and within the range of that providers certification for a particular framework.
- Introduces an approach for federated assurance by de-coupling enforcement points from decision points by adoption of a standard, open protocol.
- Postulates that there should be a fourth A added to the typical security list of AAA – and that is Assurance.

The good, the bad, and the ugly for Identity Providers

Large social sites have recently enjoyed successful "connect" business models, which have drawn many applications to them. These social sites are getting most of the business benefits from the "connect" tools by leveraging huge customer-base data to attract developers into their sites. Developers are able to use these "connect" models with standards like the OAuth 2.0 based access mechanism for getting users' "Like" data from various Graph API. For example, their social graph, relationships, likes, etc.

However, these "connect" data are not widely viewed as "trusted" source data. For example, business owners can't really rely on "Like" data for their higher-level liability/risk business.

Also, there are increasing user privacy concerns on how some sites handle privacy data. Moving from social networking services to serious eGov and eCommerce services, it is critical to develop an open standard based Trust Framework in the identity ecosystem, which is also the vision of the White House NSTIC initiative.

And there are other issues that Identity Providers will face.

Manual Onboarding

There is a need for rapid onboarding of partners through Trust Frameworks that will protect an Identity Provider from higher levels of liability and risk.

Credential/ Transactional Gap

The industry has identified a fundamental gap between credential-based trust approaches (like those used in e-gov, everything solved up front via Authentication) and transactional risk-mitigation-based

trust approaches (like those used in B2B, usually involving customized Authorization after Authentication occurs).

Overloaded Authentication

The industry is currently trying to solve all of these things by overloading the functionality of Authentication causing extremely complex technical requirement discussions by global committees that have dragged on for over 12 years.

Complex T&Cs

The Terms and Conditions for the use of a product or service are often extremely complex, with development requiring the devotion of significant time and manpower resources from legal and product teams. Even then, there is still a risk of duplication and contradiction in terms across a product portfolio, as well as a risk that important items will not be caught or effectively handled in the initial drafting. As a result, applications may be subject to potentially costly updates over time to manage any unforeseen issues.

Adaptability Costs

The ever-changing conditions in the areas of security, privacy, and legal are extremely costly for product logic adaptation in the code over time once those products are launched. Changing application code/logic to adapt also introduces new risk and "longer time to market" for new functionality.

Solution Struggles

The US Government through National Strategy on Trusted Identity in Cyberspace is currently grappling with these issues and searching for anyone that has a solution to help guide the conversation and shape the standards.

Current thinking in the industry may have gaps

New IETF standards of OAuth 2.0 and OpenID Connect provide technologies and protocols to allow identity providers to support various "connect" capabilities in an open and standard way.

Also, these are seen as having the potential to provide basic trust framework support using a Trust Claims model.

However, developing Trust Frameworks in the identity ecosystem is a huge deal, which needs new standards and industry/government collaborations like the White House NSTIC initiative. And even then few if any are approaching rapid onboarding for transactional-based assurance.

This is due to using standard approaches that are often limited to Authentication, Authorization, and Audit, which have failed to fill in some of these critical gaps.

Opportunities for an Identity/Attribute Provider (IdP) when Trust Frameworks are combined with Transactional Assurance

Increased revenue through rapid partnerships with Relying Parties.

Entities that benefit by selling goods and services to people via their Identity provided by companies like telco's, banks, large internet portals, and the like.

Increased customer base, draws people that need trusted identity to access partner sites.

Leveraging Trust Frameworks which their Identity Provider supports . These partners will include government sites, public services, health services, and more.

Increased fidelity of personal data captured in an ever-evolving Profile that each customer fills in over time to gain new functionality.

With the customer's consent this higher quality data could serve to greatly increase revenue streams from targeted advertising and other such data uses.

Introduction

In the space of Identity there is a group working to build trust in the exchange of identity credentials online called OIX (Open Identity eXchange). They set forth three major players in every online transaction: the user, the identity service provider, and the relying party. The one that seems to be consistently missing from the table of discussions are the relying parties.

In order to attract relying parties to participate, the various standards groups need to identify clear lists of features and benefits that will speak to their needs. In this case it is felt that the most important of these is providing a way for relying parties to better manage the transaction risks, and as a result gaining financial benefit. Another important benefit will be to provide the privacy-enhancing identity ecosystem with unified user consent management capability. The approach presented in this white paper addresses both of these.

Each of the three players will have an area of interest that might be more of a focus to them during an online transaction. OIX describes these areas as Levels of Control, Assurance, and Protection. The individual user may be primarily interested in their Level of Control (LOC) and in order for the identity service provider to help broker and facilitate support for that, they will need to partner with the relying party on how they measure and act upon that user's current and individual personal assurance as it may apply to the nature of the current transaction.

To that end, identity service providers are currently building platforms that contain a number of things. One of these will be an area where a person can capture their identity profile. The identity profile will allow individuals to make attribute declarations and informed agreements.

Attributes are the things that speak to our identity like our names, ages, e-mail addresses, cell phone numbers, etc.

Agreements are the contractual law binding partners for a particular discourse typically comprised of Terms and Conditions (**T&Cs)**, and various **consents** that can be applied globally to an application or service, or can apply to the specific application of that particular application or service. For example, a user can consent that the identity provided can present to them context-based advertising as an example of global consent. The user can also consent to something specific like sharing their location information with a particular person or application.

In this white paper, we postulate that the things that are considered attributes and agreements should grow significantly beyond current thinking and can be used later to speak to the assurance we have in a particular authentication.

Attributes and Agreements

Most attribute declarations should have the initial state of "not validated". They are declarations that the individual makes and may be largely unsubstantiated initially. And this is okay. There's been a great deal of effort and work lately in the industry around the area of validating identity declarations. While this work is certainly important and even critical in order for the Internet to support commerce that requires a higher level of liability, risk and responsibility, we shouldn't lose sight of the fact that aspirational declarations may actually be more valuable than proven declarations.

If a person declares to be a mountain climber do they really need to prove it to us before we charge a sporting goods company money to show that person an advertisement for rope? The person who can actually prove their declaration that they have been a mountain climber for twenty years may not need the rope as much as the person who aspires to be a mountain climber but hasn't started.

When a person first registers their identity with an identity service provider they will persist into the system a number of initial declarations and agreements. That number should initially be relatively small. Analysis has proven that the larger the number of questions there are on the initial interview the higher the abandonment rate is for any particular flow. Therefore identity systems should be able to support some functionality with initially a very small set of attributes and must be able to grow over time as needed.

Although most attribute declarations and agreements should be at first treated as un-validated, some by their very nature can be captured as validated from the very beginning. For example T&Cs may not need any additional validation beyond the individual checking a checkbox that they agree to submit to the language of a particular contract. Since the platform that supports attribute declarations and agreements will be built to support un-validated declarations, if it is ever determined that some validation beyond acceptance of T&Cs is necessary an adjustment can be made.

The same technique can apply to things like biometrics (voice passphrases and the like) that can be considered validated upon recording if a strong authentication is used to access the recording system.

Certainly for each of these a validation apparatus must be built when required. The validation apparatus may be completely electronic or a combination of electronic with human interaction. For example if an

individual declares that their username is famous.person@famousperson.com, we may be able to trust that declaration to a certain degree but at some point we will need that declaration to be validated to prove that they indeed have access to *that* famous person email account. In this case the validation apparatus may be that the system sends an e-mail to *that* e-mail address and asks the user to login to their e-mail client with the correct username and password combination and click on a hyperlink or enter a PIN into a user-interface from that e-mail. Once received and processed that round-trip will prove that the person who has access to that e-mail address has taken the action to validate that attribute declaration.

For all validations there is a **date of creation**. But we need to go beyond that. When a person validates a declaration we can capture a **date of establishment**. But we probably don't want to trust that data as valid forever. Therefore each identity service provider must work with their own security, privacy, and legal advocates to determine the rules for how long they can trust data once it has been established in the system. These rules may be unique to each attribute type and agreement.

Based upon these rules for each attribute type the system will be able to determine a **date of suspect**, or a date when we can no longer trust that validation. And yet just because we no longer trust that data doesn't mean that we can immediately release it to somebody else. For example if a person purchases a cell phone and registers that new telephone number as an attribute declaration on their identity and then subsequently goes through a validation apparatus to have a PIN sent to that cell phone creating a date of establishment, the issuing company will know immediately when that data can no longer be trusted-it's the date that the contract expires. Therefore the date of suspect can be set to the date the contract ends. But by federal guidelines that telephone number cannot immediately be released for somebody else to use when the contract expires. It must be set aside for 90 days before it can be released. So the system that supports attribute declarations must be able to set dates of creation, establishment, suspect, and **date of release** for probably each attribute type.

Platform Considerations

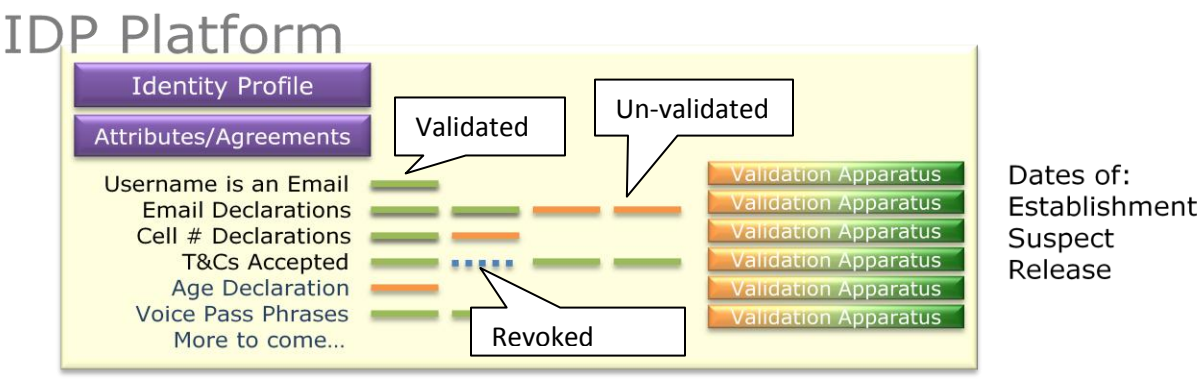

Figure 1

Figure 1 is a graphical representation of data that a user might enter into a system with mixed states of validated, un-validated, and revoked. Each attribute and agreement type has a validation apparatus, and each datum entered has a date of establishment, suspect, and release.

For all attributes and agreements the person must be given the opportunity for revocation of each declaration individually. They may cancel a cell phone, change their last name, or no longer consider themselves affiliated with a particular group. This would include T&Cs or any other attribute declaration – even biometrics.

This doesn't mean that the user would escape contractual executory responsibility. For example they would still have to pay their monthly cell phone bill, but they could say that they no longer wish their cell number to be associated as an attribute of their identity, or they no longer agree with a specific T&C.

Once the system has been built to support identity profiles with attribute declarations, two more components need to be added - a **decision point** and a **remedy point**.

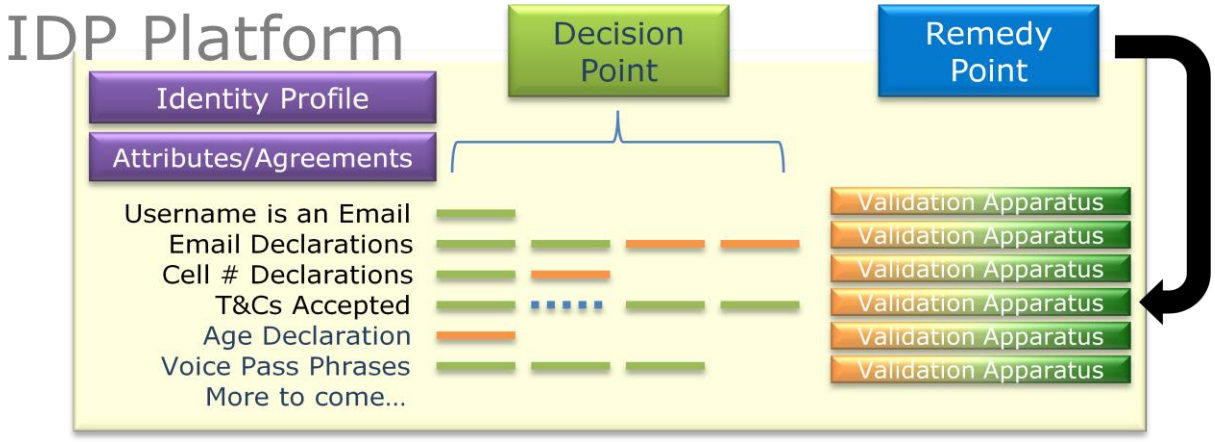

Figure 2

Decision Point

The decision point is a mechanism that can look at the overall landscape of current attribute declarations/agreements and their various states of validation and answer two kinds of questions.

Decision Point Question Type #1 - PLOA

The first is a true/false reply when presented with a set of conditions. The set of conditions should support hierarchical questions. For example the model supports the following types of questions:

- Does the user have *any* declarations that have been validated?
- Does the user have any validated *e-mail address*?
- Has the user validated their *primary* e-mail address?
- Has the user validated this *specific string* that is an e-mail address?

It should also support combinations of conjunctions like does the user have any e-mail address validated *and* have they accepted a specific T&C?

If the response is false to any request a standard error that is appropriate for that transport and technology is generated with the necessary information for the calling app to respond in a standard way. For example in HTTP, a 403 error is generated by the system and the XML body of that error will contain a standard code and enough minimal description to facilitate all the validations necessary to elicit a desired true response without divulging private or sensitive information about the user.

When the calling app is created by a third party and privacy is a concern, a mechanism can be put in place to obfuscate the actual reason for the error. For example a unique one-use code can be generated by the system and passed in the body of the error message to the calling app while simultaneously being communicated to the appropriate remedy point. We will cover **remedy points** in detail later, but suffice it to say that it is where a UI can go to resolve some issue with the current state of attribute validation or list of agreements that have been accepted. The UI can call the Remedy point and present the code and be directed to take the user through the correct steps to validate the necessary attribute declarations or make the additional agreements. Conversely the Decision point could broker that conversation for the UI orchestrating to the Remedy point and passing back the necessary components to the UI directly.

We will cover the subject of **enforcement points** in detail later, but suffice it to say that they form the security perimeter around protected services. The main thing to remember here from a security/privacy point of view is that the enforcement point has been granted access to the decision point only via an established trust relationship between the provider and the relying party. Also there is a presumption that a valid and secure authentication mechanism is in place. Lastly that the UI application can be limited in access to a remedy point via an established trust relationship if the data being shared is ever considered sensitive and can be correctly managed and audited over time.

Decision Point Question Type #2 - LOA

The second question is an overall assessment of the current state of the attribute data for a specific user boiled down to a single number that provides assurance based upon the rules set forth by the security, privacy, and legal advocates that manage that particular IDM platform. In these cases agreements such as T&Cs or various consents are typically less important to relying parties.

For example if a particular identity service provider has been certified to be a level 3 by some standards organization and can speak to the general level of assurance (or LOA) of an individual, the decision point should be able to look at the overall landscape of an individual's attribute declarations and determine if that user should be considered a 1, 2, or 3 per their own internal approach which was certified to do so. Companies will be free to develop their own secret sauce for how they weight attribute validation, will still need to get that secret sauce to be certified per some auditing mechanism to support a certain standard, and yet share that information in a common and simple format. This way the approach engenders competition and innovation while exposing consistent interfaces and common data formations.

Cross Platform/Partner Protocol

Therefore this white paper proposes that decision points shall be standardized to respond with either a 1) true/false or 2) a level in such a way that supports truly federated communications. This suggests the need for the adoption of an agreed-upon common protocol.

Remedy Point

The remedy point is a standard service that simply redirects to the appropriate validation apparatus to resolve the current issue. If it's built right the remedy point will be able to respond to user interfaces in the appropriate way for that particular technology. For example if called by a Flex/Flash client the remedy point can respond by serving up a .swf file with the necessary components that will allow the user to validate a certain combination of declarations. But if the remedy point is called by a client that is HTML 5 running in a browser it could provide a browser redirect to an HTML 5 page where the user could perform the validation with a subsequent call-back returning the user back to their application.

In most of our analysis to date the typical validation requires some input from the user. In the case of a cell phone validation the user will likely need to enter a PIN that is sent to their declared phone number. For a T&C the user would have to be able to read the language and agree to it. For a voice print or other bio-metric they would have to be able to capture the impression and submit it. If constructed correctly the UX for this can be a dynamically rendered combination of sub-modules that can help the user resolve any/all necessary validation in one step. However, it is likely that the remedy may require more actions than can be accommodated by the present UX. In these cases the appropriate text is presented so the user knows what to do and control is passed back to the application. The user may act upon the same protected object again and simply get the same results until they remedy the necessary validation.

Because the decision point and remedy point are called at the moment a verification of previous validations is needed, the user's ability and right to revoke any or all attribute declarations at any time can be supported. For example if a user agrees to T&C #14 in January and then decides to revoke that declaration in March the decision point and remedy point if called in June can return a false and redirect the user to once again agree to the T&Cs necessary to support a particular functionality.

Data Center Application Deployment

Now that we have covered some basics regarding identity profile and attribute declarations and the necessary infrastructure to support validation and subsequent verifications, let's talk about an application that has been deployed to a data center.

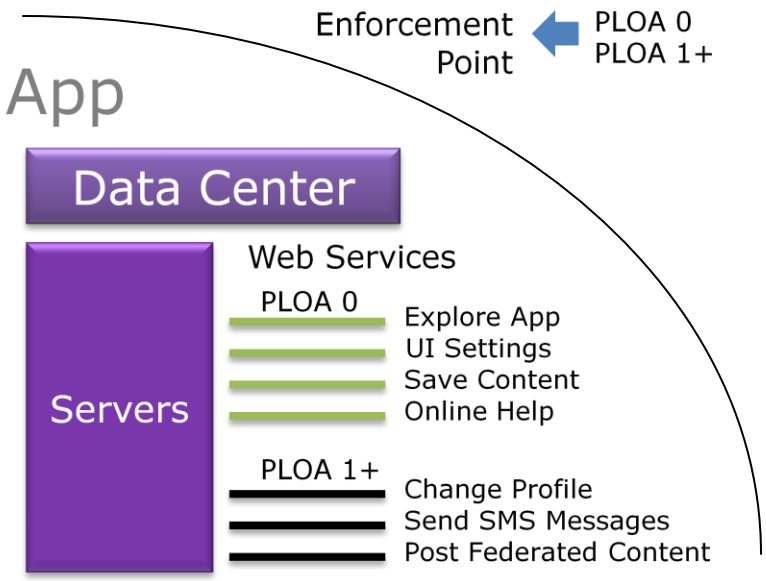

Figure 3

All application data centers have certain elements in common. All data centers have servers where applications have been installed. They all have an enforcement point that can be as simple and inexpensive as firewalls and load balancers or as robust as expensive devices that can handle XML acceleration and deep packet inspection and can be taught rules. What is it that the enforcement point is trying to protect? API running on the servers! For the purposes of this white paper we will use Web services as an example of such APIs.

PLOA of 0 (zero)

These Web services can be determined as posing a very low level of liability to the company or the user. For example no one will sue a company because they change their background color from blue to purple. So Web services that support things like exploring an application, changing your UI settings, reading T&Cs, or accessing online help might be deemed by that company's legal, privacy, and security advocates as posing a very low liability. But some Web services such as changing your profile, or sending SMS messages which incur a charge on your phone bill, or posting federated content to a partner website might be considered to pose a higher level of liability to the company or the individual.

In the case where there is a very low level of determined liability we might establish a PLOA of zero or having no assurance verified higher than plain old authentication.

PLOA of 1+

But if the Web services that we are protecting pose a higher level of liability or risk exposure then they should be protected with a PLOA of 1 or higher – meaning that an attribute or some combination of attribute declaration has been previously validated by the user and the present date falls between the date of establishment and the date of suspect.

These rules apply to atomic API. Each separate Web service that may be exposed can be taught to the robust enforcement points as PLOA zero or higher. Each web service can easily have unique conditions that must be in place before the company would feel comfortable serving up the content to or respond to the request of an authenticated user. In this regard PLOA would be providing personal assurance at the point that the web service is being called and would speak to the assurance level that company has on that authentication.

The Four A's

Even though the technology being implemented may resemble authorization, it is truly speaking to the assurance of the authentication and therefore should be considered a new element to the three A's.

Authentication, Authorization, and Audit (AAA) shall be joined by their new sibling Assurance. So we should now consider that there are four A's in the area of cyber security.

Low Liability Call Flow

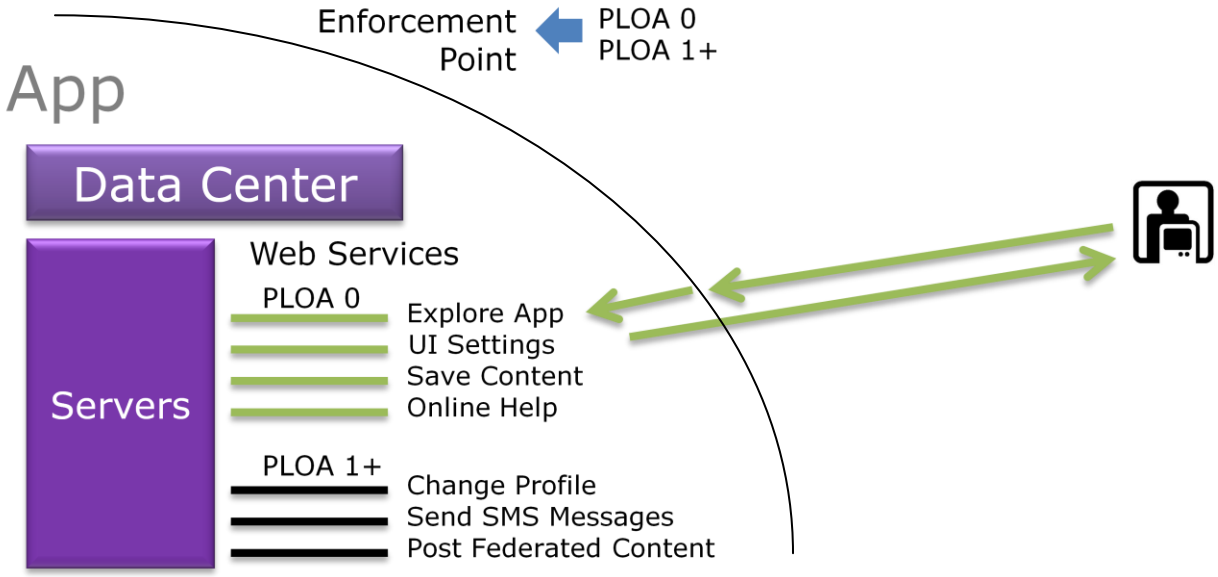

Figure 4

If the user desires to access Web services that have a very low liability or a PLOA of 0 (zero) then their user interface would call those Web services directly and it would be intercepted by the enforcement point. The enforcement point would examine the web service in light of the rule that it was taught, determine that only authentication is required, would allow that transaction to continue unmolested, and the user would receive back the desired content.

What would happen if the user desired to access Web services that were thought to have a very high liability or a PLOA of 1 (one) or higher?

High Liability Call Flow

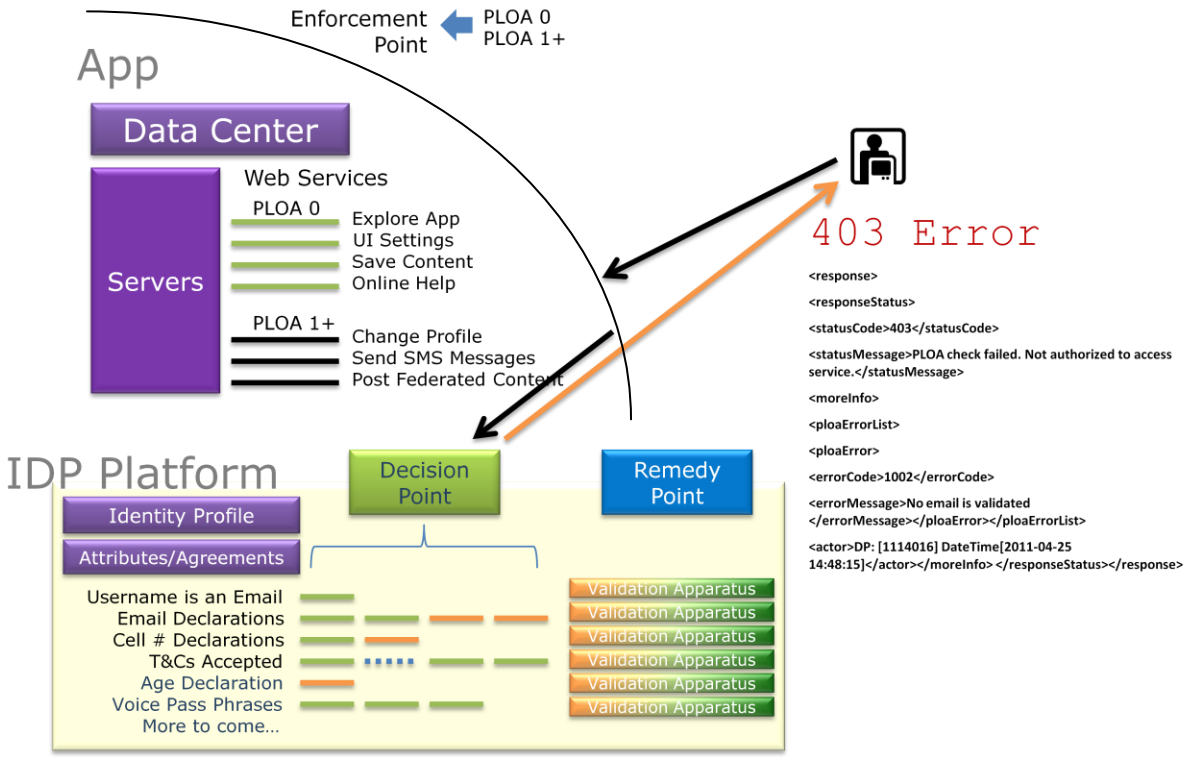

Figure 5

The UI calls that specific Web service and that call is intercepted by the enforcement point. The rule protecting that web service causes the enforcement point to call the decision point and presents to it all the data that it has in that moment. The decision point then replies with either a true, or a false that is a 403 error with the necessary code and language for the UI to understand how to remedy the situation while obfuscating the actual details as mentioned previously.

If a false is returned to the enforcement point the enforcement point would then proxy that information back to the requesting UI. The UI could then call the remedy point and present to it the error code and information and the remedy point would redirect the UI to the correct validation apparatus so that the user could validate the necessary information to support that particular call.

The user's experience in this particular exercise would be something like this:

The user clicks on a button to request sensitive information, the user is presented with a new T&C, the user accepts the T&C and is presented with the desired content. In the future when that user clicks on the same button they are simply presented with the desired content and no request to agree to a T&C. But if conditions change for the company and additional validation is required based upon new company policy and/or legal direction the user might see more validation steps for one time upon request. In this way the company has the flexibility to alter the conditions necessary to protect itself without the need

to change code on the server or the client side. And the user understands with clarity the changing conditions that have occurred and can either agree to them or not.

In-transaction Identity Proofing

According to the NIST 800-63 document, there are two major LOA activities related to the user, Identity Assurance and Authentication Assurance.

Identity Assurance is normally done by Identity Proofing *before* the identity can be vetted. For example, the user needs to present two photo ID documents to a Registrar to validate the user identity registration data for LOA 3, etc.

Normally, Identity proofing is done after user identity registration before the user can declare more attributes to his/her identity profile.

But by using the above process, PLOA is suggesting a late binding method that can be used to combine Identity Proofing into the Attribute Validation process.

Modular T&Cs

The use of this method to develop a modular approach to T&Cs could greatly simplify the current process, leading to flexibility and cost savings for all parties. Rather than a single, comprehensive set of terms that must be developed by service providers and agreed to by users prior to implementation, covering a variety of circumstances that might never apply in that particular relationship, the parties instead could rely upon a modular approach in which terms are developed and agreed to on an incremental basis over time as previously unforeseen items arise. By applying to very atomic functionality that can be exposed across multiple applications, this approach could potentially avoid costly after-the-fact logic changes in the application's code necessitated by unforeseen changes in circumstance. For example a single T&C could be written to support a function like SMS messaging from a browser that, once agreed, to could cover this functionality exposed in many different apps. Such an approach could also work to the benefit of both service providers and users by simplifying and making the overall process more easily understandable on both sides of the table.

Further, users will not have to validate *any* attribute declaration until it is absolutely necessary for specific atomic functionality that exposes the company to higher liability. This can greatly decrease the number of attributes captured at initial registration which in turn will greatly reduce the abandonment statistics to any product flow. More users will use more platforms with confidence to determine whether they are a good fit for their lifestyle – which can only be good news for product developers.

Federated Assurance

Now let's examine how this may play out in a federated relationship where one company is an identity service provider with their own enforcement point, decision point, and remedy point. And the other company is a relying party with their own data center, Web service, remedy point, and decision point.

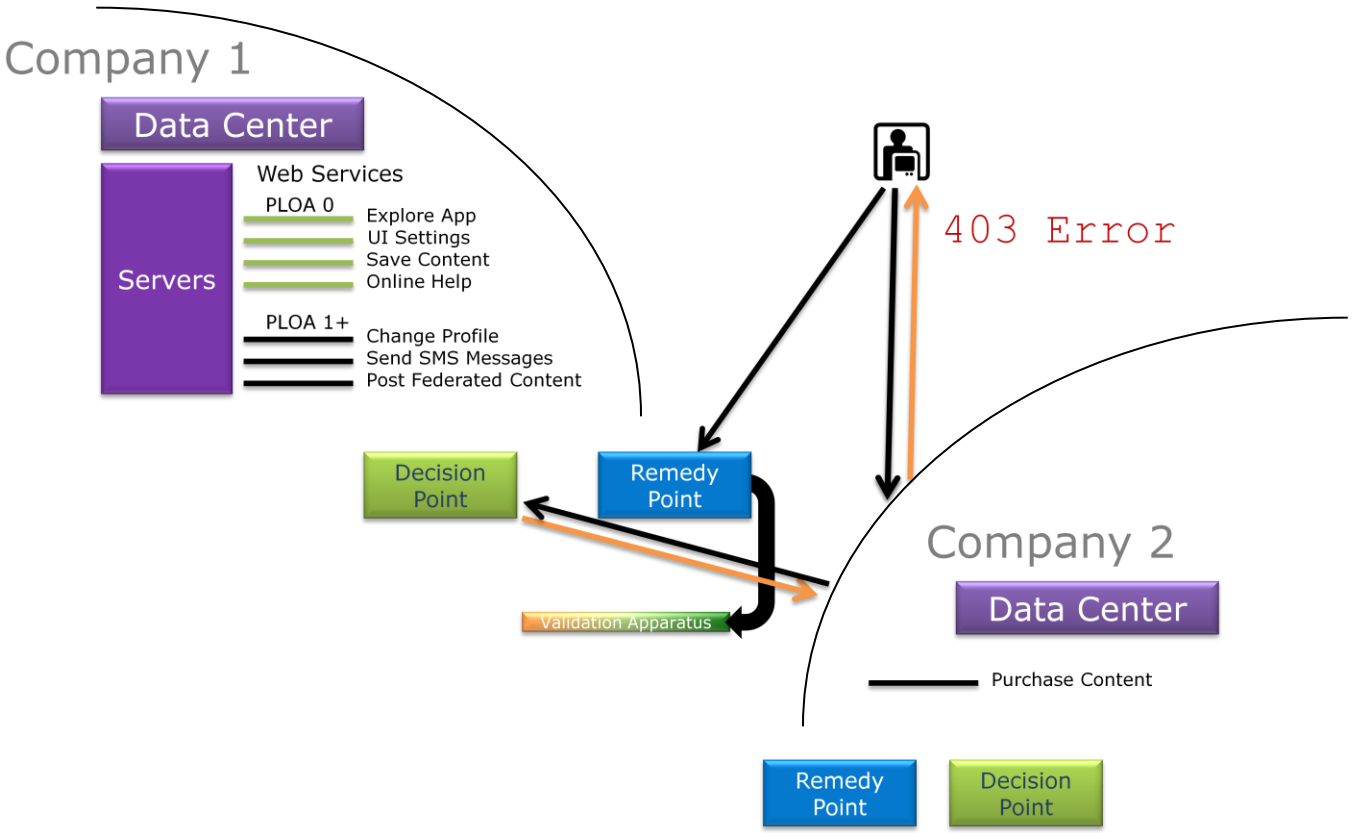

Figure 6

The user running a third-party application downloaded from an application store has authenticated with the identity service provider and has a token in the header provided via that authentication. The user desires to call the web service that resides in the relying party's data center. The transaction is intercepted by the relying party's enforcement point which being federated with the identity service provider's decision point calls that decision point and presents all the data necessary to support the relying party's rules for PLOA that were taught to the relying party's enforcement point. The identity service provider's decision point can return a 403 error which is proxied back to the third party's UI. In turn the UI would call the identity service provider's remedy point because it has their token in its header. The remedy is presented to the user and the user takes action to validate the necessary attribute declarations. Now the UI tries again and when the relying party's enforcement point calls the identity service provider's decision point it gets back a true and allows the transaction to continue inside their own firewalls unmolested and the user gets back the desired content.

Real Privacy

If the remedy point is limited via a high-level relationship with a broad base of relying parties, and as such limits access to only general questions and not specific values, then the relying party only knows that generic questions like, "Johnsmith@mail.com and valid primary email address" is currently true – and that may be enough to satisfy their lawyers that they have safety from liability or risk.

In this way, the identity service provider never gains knowledge of what the user is doing. It just knows that a trusted partner requested a verification of validated declarations and responded in kind.

Conversely the relying party can be ignorant to the actual values of the attributes. In the example their minimal liability requirements have been met via validation by a partner IDP per a business relationship with them, and that would be sufficient to mitigate some perceived liability exposed by a particular functionality served up by their company. In this case their enforcement point may simply state “this user and validated primary email address” to the IDP’s decision point. They could get back a “true” and never actually need to know that customer’s email address.

Ultimately this speaks to the customer’s Level of Control, and provides an architecture where the IDP is ignorant to the customer’s detailed activity and the relying party is ignorant to the customer’s attribute values. Therefore, PLOA goes beyond providing a heightened assurance to a person’s authentication and actually provides functionality to support the customer’s Level of Control.

This underlines the earlier proposition that all decision points shall reply in a standard way that is ultimately a protocol.

Performance Considerations

We realize that performance will be a strong requirement to get right since this will be in-path of live network requests. The protocol must be kept extremely lightweight. Also, the method by which the decision point caches data should be examined. There are developing architectures that securely support caching of data on the edge of an implemented stack via clustering and highly scalable data distribution.

IDPs should realize that even though “how” they implement their stack is completely up to them, slow performance of any particular IDP will likely have a negative impact on their own user adoption. It will be in every IDP’s best interest to make their edge for their decision point to be as responsive as possible. In this way this approach will support one of the main goals for the National Strategy for Trusted Identities in Cyberspace (NSTIC) – competition in the area of privacy and identity – while at the same time providing minimal common infrastructure that can be scalable and easier to onboard partners.

PLOA support for LOA Value Setting

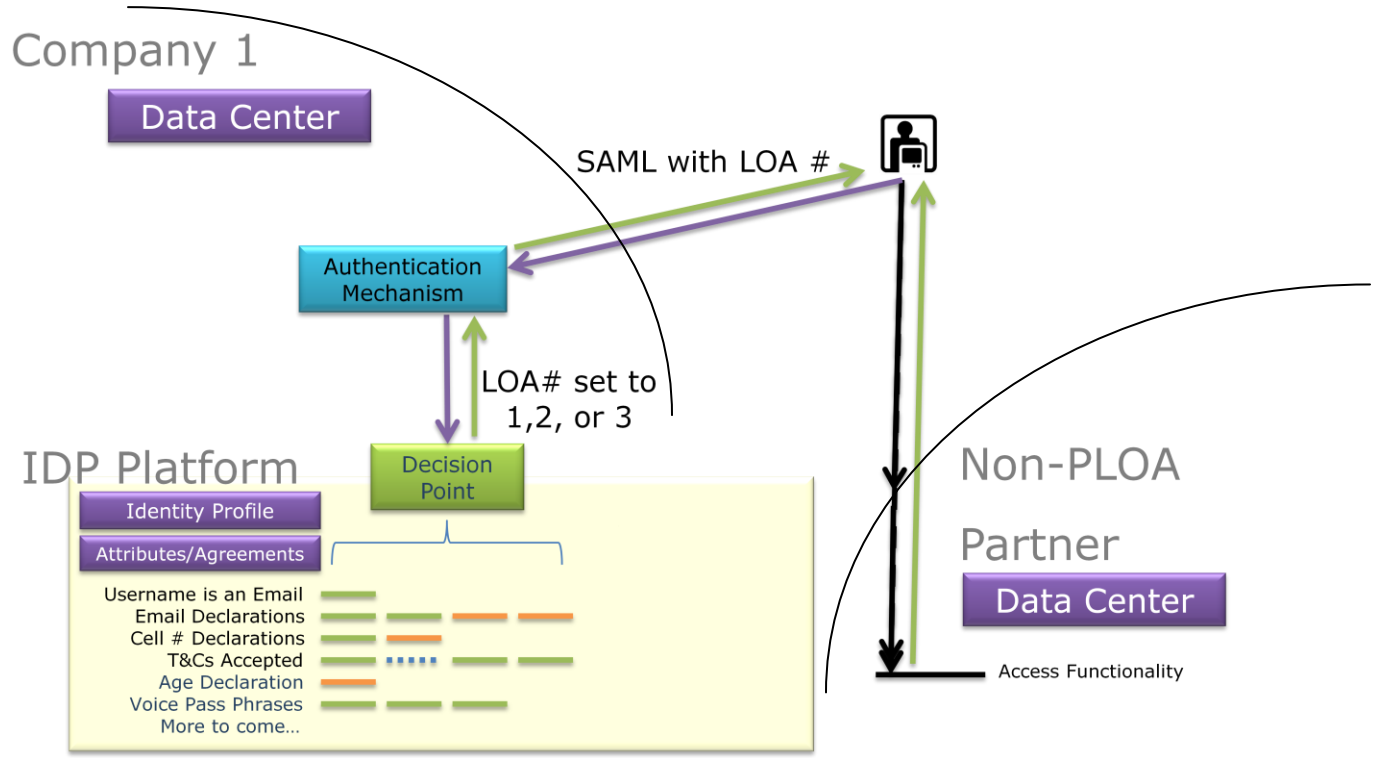

Figure 7

Not all companies will be able to afford robust enforcement points that can be taught rules for individual PLOA assigned to specific Web services. In these cases a more boiled down approach to traditional Level of Assurance (LOA) needs to be supported. Given the above precepts let's examine what would happen in the case where a relying party is one such company.

One company is an identity service provider with a data center, an enforcement point, decision point and remedy point. The other is a service provider that is a relying party with its own data center, web service, but doesn't have a remedy point, a decision point, or any PLOA support.

The user running a third-party application downloaded from an application store has authenticated with the identity service provider. The identity service provider's authentication mechanism called their own decision point and requested an assessment of the overall landscape and current conditions of attribute declarations based upon their own algorithm, and decided that this particular user is a 1, 2 or a 3 per their certification as a company that has an LOA of 3. The identity service provider sets the specific bit in the generated SAML assertion (or other agreed-upon federated communication mechanism) to that specific LOA for that user at that moment. Now when the third-party UI calls the non-PLOA partner, their enforcement point simply looks at the SAML assertion and determines that the LOA is appropriate for the content that they are serving and responds accordingly per establish standards and guidelines.

In this way the decision point of the identity service provider that hosts the attribute declarations can either respond to atomic requests for very specific conditions via PLOA or can respond to an overall single number in order to establish a federated authentication to a specific LOA.

Conclusion

By offloading the burden of Assurance from various Authentication mechanisms, we greatly decrease the complexity around Authentication federation where various approaches have faced complex and even stifling obstacles when dealing with technical implementation of overloaded and even bloated Authentication responsibilities.

This white paper proposes that by combining both PLOA and LOA assessment functionalities into a single decision point used for in-line query of a specific user's assurance at a particular moment in time (either post-authentication or at the time of federated authentication), which can also be federated with one or more enforcement points, new flexibility can be introduced to the development and consumption of products and services across multiple industries and will ultimately safeguard the privacy, protection and rights of the individual.

Recommendations

- The industry should collaborate immediately on an Attribute and Agreement Taxonomy. Standard classifications for each type of attribute/agreement, and the comparable and acceptable rules for each provider to define Dates of Suspect/Release need to be established as guidelines.
- A protocol decision needs to be made. For PLOA in regards to consistently federating an enforcement point with a decision point, should there be a new protocol or should an existing one be adopted? What are the governing requirements for that protocol?

Glossary of Terms

Term	Definition
AAAA	Authentication, Authorization, Audit, and Assurance
ASP	Attribute Service Provider - A company set up to allow users to create, read, update, and delete attributes associated with a particular identity and to some degree federate them with consent.
Attribute Declaration	The moment a person initially persists an Attribute data into their Profile, and typically it is considered un-validated.
Attribute Revocation	The moment a person takes action to disassociate a particular attribute from their identity.
Attribute Validation	The process that results in a user proving that a certain attribute declaration is true.
Attributes	Some element of data that can further define an identity – like "my first name is John", John would an attribute.
Decision Point	A mechanism that can look at the overall landscape of current attribute declarations and their various states of validation and reply with a True/False or an overall evaluation conforming to some standard.
DoC	Date of Creation – the timestamp an Attribute is first declared.
DoE	Date of Establishment – the timestamp when a user successfully completes validation of an Attribute Declaration.
DoR	Date of Release – the timestamp the Attribute Provider can allow another individual to lay declare ownership of a particular attribute data. This may not universally apply as some attributes can be shared, but does work well for those that are uniquely owned.
DoS	Date of Suspect – the timestamp the Attribute Service Provider can no longer trust an attribute's validation based upon their own rules and guidelines.
IdP (IDP)	Identity Service Providers – typically companies that can allow a person to create an "identity", some method the authenticate it (username/password), and to some degree federate it.
LOA	Level of Assurance – typically an authority's appraisal of the overall assurance that an IDP can offer them based on results of an audit. Often LOA is at the center of technical standards for sharing mechanisms as well as types of Authentication.
LOC	Level of Control
LOP	Level of Protection
OIX	Open Identity eXchange - working to build trust in the exchange of identity credentials online.
PLOA	Personal Levels of Assurance
Remedy Point	The remedy point is a standard service that simply redirects to the appropriate validation apparatus to resolve a current issue.
T&C	Terms and Conditions – a legal contract users agree to in order to gain access to some application or functionality.
Validation Apparatus	Sometimes called Attribute Proofing, it is the mechanism (automated and/or manual) that facilitates Attribute Validation for a particular attribute type.